



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office

Address: COMMISSIONER FOR PATENTS

P.O. Box 1450

Alexandria, Virginia 22313-1450

www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/760,659	01/20/2004	Simon C. Steely JR.	200313629-1	9866
22879	7590	04/29/2010		
HEWLETT-PACKARD COMPANY			EXAMINER	
Intellectual Property Administration			CHERY, MARDOCHIEE	
3404 E. Harmony Road				
Mail Stop 35			ART UNIT	
FORT COLLINS, CO 80528			PAPER NUMBER	
			2186	
			NOTIFICATION DATE	
			DELIVERY MODE	
			04/29/2010	
			ELECTRONIC	

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

JERRY.SHORMA@HP.COM

ipa.mail@hp.com

laura.m.clark@hp.com



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

**BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES**

Application Number: 10/760,659
Filing Date: January 20, 2004
Appellant(s): STEELY ET AL.

Steven L. Nichols
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed February 10, 2010 appealing from the Office action mailed December 10, 2009.

(1) Real Party in Interest

The examiner has no comment on the statement, or lack of statement, identifying by name the real party in interest in the brief.

(2) Related Appeals and Interferences

The following are the related appeals, interferences, and judicial proceedings known to the examiner which may be related to, directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal:

10/760,652.

(3) Status of Claims

The following is a list of claims that are rejected and pending in the application:

1, 8-12, 18-32.

(4) Status of Amendments After Final

The examiner has no comment on the appellant's statement of the status of amendments after final rejection contained in the brief.

(5) Summary of Claimed Subject Matter

The examiner has no comment on the summary of claimed subject matter contained in the brief.

(6) Grounds of Rejection to be Reviewed on Appeal

The examiner has no comment on the appellant's statement of the grounds of rejection to be reviewed on appeal. Every ground of rejection set forth in the Office action from which the appeal is taken (as modified by any advisory actions) is being maintained by the examiner except for the grounds of rejection (if any) listed under the

subheading "WITHDRAWN REJECTIONS." New grounds of rejection (if any) are provided under the subheading "NEW GROUNDS OF REJECTION."

(7) Claims Appendix

The examiner has no comment on the copy of the appealed claims contained in the Appendix to the appellant's brief.

(8) Evidence Relied Upon

20020129211	Arimilli	9-2002
6,931,496	Chen	4-2003
6,484,240	Cypher	11-2002

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 1, 8-10, 12, 18-19, and 24-32 are rejected under 35 U.S.C. 103(a) as being unpatentable over Arimilli (2002/0129211) in view of Chen (6,931,496).

As per claim 1, Arimilli discloses a system comprising: a first node including data having an associated state, the associated state of the data at the first node being a modified state [page 7, par. 0060, claim text 1]; and a second node operative to provide a non-migratory source broadcast request for the data, the second node being operative to receive the data from the first node and assign a shared state to an associated state of the data at the second node [par. 0029, page 7, par. 0060, 0062, claim text 1].

However, Arimilli does not explicitly teach the first node being operative in response to the non-migratory source broadcast request to provide the data to the second node and to transition the associated state of the data at the first node from the modified state to an owner state without updating memory as required by the claim.

Chen discloses the first node being operative in response to the non-migratory source broadcast request to provide the data to the second node and to transition the associated state of the data at the first node from the modified state to an owner state without updating memory [col. 5, ll 5-10, ll 36-42, ll 60 to col. 6, ll 6] to provide a data maintenance method in a distributed shared memory system to efficiently solve the access deadlock problem (col. 3, ll 10-12).

Since the technology for implementing a cache coherency system with transitioning the associated state of the data at the first node from the modified state to

an owner state without updating memory was well known as evidenced by Chen, an artisan would have been motivated to implement this feature in the system of Arimilli to provide a data maintenance method in a distributed shared memory system to efficiently solve the access deadlock problem. Thus, it would have obvious to one of ordinary skill in the art at the time of invention by Applicant, to modify the system of Arimilli to include transitioning the associated state of the data at the first node from the modified state to an owner state without updating memory because this would have provided a data maintenance method in a distributed shared memory system to efficiently solve the access deadlock problem (col. 3, ll 10-12) as taught by Chen.

As per claim 8, Chen discloses the first node is operative in response to the non-migratory source broadcast request to provide a shared data response to the second node [col. 6, ll 1-6].

As per claim 9, Arimilli discloses further migration of the data from the second node is precluded when the associated state of the data at the second node is the shared state [Abstract].

As per claim 10, Chen discloses at least one other node that provides a non-data response to the second node in response to the non-migratory source broadcast request from the second node, the non-data response indicating that the at least one other node does not have a valid copy of the data requested by the second node [col. 7,

II 55-61].

As per claim 12, the rationale in the rejection of claim 1 is herein incorporated. Arimilli further discloses second processor nodes having cache lines with associated states, and that the second processor node is operative to provide a non-migratory read request for data [Fig. 1; par. 0029].

As per claim 18, the rationale in the rejection of claim 9 is herein incorporated.

As per claim 19, the rationale in the rejection of claim 10 is herein incorporated.

As per claim 24, the rationale in the rejection of claim 1 is herein incorporated.

As per claim 25, the rationale in the rejection of claim 1 is herein incorporated. Arimilli further discloses means for providing the data from the second node to the first node in response to a migratory read request [Figs. 3-4; *master receives store operation from processor, owned state*; par. 0029]; and means for transitioning the modified state associated with the data at the second node to an invalid state [Figs. 3-4, *master performs store, issues kill transactions to invalidate copies held by other agents*; pars. 0006-0008].

As per claim 26, Chen discloses means for selecting one of the XREADM request and XREADN request to broadcast from the first node [Figs. 4 and 5].

As per claim 27, Chen discloses means for predictively selecting one of the XREADM request and XREADN request to broadcast from the first node [Figs. 4 and 5].

As per claim 28, the rationale in the rejection of claim 1 is herein incorporated.

As per claim 29, the rationale in the rejection of claim 1 is herein incorporated. Arimilli further discloses broadcasting from a first node a migratory read request for data responsive to the migratory request for the data [Figs. 3-4; par. 0029]

As per claim 30, the rationale in the rejection of claim 26 is herein incorporated.

As per claim 31, the rationale in the rejection of claim 27 is herein incorporated.

As per claim 32, the rationale in the rejection of claim 1 is herein incorporated. Arimilli further discloses a cache coherency that permits migration of data to a cache associated with a source processor from a cache associated with a target processor when a migratory request is issued [Figs. 3-4; *master receives store operation from processor, owned state; master performs store, issues kill transactions to invalidate copies held by other agents; par. 0029*].

Claims 11 and 20-23 are rejected under 35 U.S.C. 103(a) as being unpatentable over Arimilli (2002/0129211) in view of Chen (6,931,496) and further in view of Cypher (6,484,240).

As per claim 11, Arimilli discloses the first node defines a first processor and the second node defines a second processor [Fig. 1], the first and second processors each having an associated cache that comprises a plurality of cache lines [Fig. 1], the first and second processors being capable of communicating with each other and with a system memory via an interconnect [Fig. 1, System Bus 12].

Chen further discloses the system further comprising a first cache controller associated with the first processor and a second cache controller associated with the second processor [Fig. 1], the first cache controller being operative to manage data requests and responses for the associated cache of the first processor [Fig. 1], the first cache controller effecting state transitions associated with the data in the associated cache of the first processor based on the data requests and responses for the associated cache of the first processor [Fig. 1], the second cache controller being operative to manage data requests and responses for the associated cache of the second processor [Fig. 1], the second cache controller effecting state transitions associated with the data in the associated cache of the second processor based on the data requests and responses for the associated cache of the second processor [Figs. 2-5].

However, Arimilli and Chen do not explicitly teach each cache line having a respective tag address that identifies associated data and each cache line having state information that indicates a state of the associated data for the respective cache line as required by the claim.

Cypher discloses each cache line having a respective tag address that identifies associated data and each cache line having state information that indicates a state of the associated data for the respective cache line [col. 1, ll 20-25] to specify the access rights and ownership responsibilities for a corresponding processor (col. 1, ll 23-25).

Since the technology for implementing a cache coherency system with each cache line having a respective tag address that identifies associated data and each cache line having state information that indicates a state of the associated data for the respective cache line was well known as evidenced by Cypher, an artisan would have been motivated to implement this feature in the system of Arimilli and Chen in order to specify the access rights and ownership responsibilities for a corresponding processor. Thus, it would have been obvious to one of ordinary skill in the art at the time of invention by Applicant to modify the system of Arimilli and Chen to include each cache line having a respective tag address that identifies associated data and each cache line having state information that indicates a state of the associated data for the respective cache line since this would have enabled specifying the access rights and ownership

responsibilities for a corresponding processor (col. 1, ll 23-25) as taught by Cypher.

As per claim 20, the rationale in the rejection of claim 11 is herein incorporated.

As per claim 21, the rationale in the rejection of claims 1 and 11 is herein incorporated. Arimilli further discloses a source processor having an associated source processor cache, the source processor being operative to issues a selected one of a non-migratory source broadcast (XREADN) request for data and a migratory source broadcast (XREADM) request for the data [Figs. 3-4; par. 0029].

As per claim 22, Chen discloses the source processor further comprises an associated source processor cache having a source processor cache line for storing the data, the source processor cache line having an associated state, the source processor storing the data in the source processor cache line and assigning a shared state to the associated state of the source processor cache line in response to receiving the S-DATA response from the target processor [coll6, ll 58 to col. 7, ll 17; col. 7, ll 49 to col. 8, ll 3].

As per claim 23, Chen discloses the source processor further comprises an associated source processor cache having a source processor cache line for storing the data, the source processor cache line having an associated state, the source processor storing the data in the source processor cache line and assigning a dirty state to the

associated state of the source processor cache line in response to receiving the D-DATA response from the target processor [col. 7, ll 49-61; col. 8, ll 52-67].

Allowable Subject Matter

Claims 2-7, 13-17 are objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims.

(10) Response to Argument

In view of the following discussion, Examiner would like to emphasize the following:

Sources of rationale supporting a rejection under 35 U.S.C. 103 may be in a reference, or reasoned from common knowledge in the art, scientific principles, art recognized equivalents, or legal precedent. The CCPA has held that "in considering the disclosure of a reference, it is proper to take into account not only specific teachings of the reference but also the inferences which one skilled in the art would reasonably be expected to draw therefrom." In re Preda, 401 F.2d 825, 826, 159 USPQ 342, 344 (CCPA 1968); MPEP 2144.01.

In determining obviousness under 35 U.S.C. 103 in view of the Supreme Court decision in KSR International Co. v. Teleflex Inc., the Supreme Court stated that: "If a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill".

Still further, the Court states that "the focus when making a determination of obviousness should be on what a person of ordinary skill in the pertinent art would have known at the time of the invention...and this is regardless of whether the source of that knowledge and ability was documentary prior art, general knowledge in the art, or common sense".

In light of the forgoing, Examiner would like to accentuate that the combinations of Arimilli and Chen is relied upon in rejecting claims 1, 8-10, 12, 18-19, and 24-32, and the combinations of Arimilli, Chen, and Cypher in rejecting claims 11 and 20-23. In response to applicant's arguments against the references individually, one cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986).

Appellant's arguments on pages 13-14 of the remarks that Arimilli fails to teach or suggest that "the second node is operative to receive data from the first node and that the second node is operative to assign a shared state to an associated state of the data at the second node", recited in claim 1, are clearly erroneous.

Examiner respectfully disagrees with Appellant's assumption. Arimilli clearly discloses "master 26 of a first agent, for example, processor complex 10a (i.e., first node), issues modifying transaction 150a on system bus 12 that targets a cache line

that is indicated as shared in the cache directory 22 of at least one agent 10 (i.e., second node); a result of a store request by the processor 16 in the second agent 10 that targets a cache line marked as shared in the second agent's cache processor 22;
par. [0029]."

In response to applicant's arguments against the references individually, one cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986).

Chen also can be relied upon for such teaching as shown in the following: "at least a first and a second nodes; the first node includes an external cache for storing a data from a local memory of the second node and at least two processors optionally accessing the data from the external cache; whether a second certain one of at least two processors is allowed to share the modified data is further determined; if the second certain processor is allowed to share the modified data, it may directly request the modified data from the first certain processor; Abstract".

Furthermore, Examiner would like to point out that the functional recitations of "the second node is [capable of or] operative to receive data from the first node and that the second node is [capable of or] operative to assign a shared state do not make the claimed invention patentably distinct over the prior art of record. While features of a system may be recited either structurally or functionally, claims directed to a system must be distinguished from the prior art in terms of structure rather than function. *In re*

Schreiber, 128 F.3d 1473, 1477-78, 44 USPQ2d 1429, 1431-32 (Fed. Cir. 1997) (The absence of a disclosure in a prior art reference relating to function did not defeat the Board's finding of anticipation of claimed apparatus because the limitations at issue were found to be inherent in the prior art reference); see also *In re Swinehart*, 439 F.2d 210, 212-13, 169 USPQ 226, 228-29 (CCPA 1971); < *In re Danly*, 263 F.2d 844, 847, 120 USPQ 528, 531 (CCPA 1959). "System claims cover what a device *is*, not what a device *does*." *Hewlett-Packard Co. v. Bausch & Lomb Inc.*, 909 F.2d 1464, 1469, 15 USPQ2d 1525, 1528 (Fed. Cir. 1990) (emphasis in original).

Still further, claim 1 recites, inter alias, "the second node is operative to receive data from the first node and that the second node is operative to assign a shared state to an associated state of the data at the second node". While these limitations require "a second node" and "a first node", the functional recitations: "receiving data" and "assigning a shared state to an associated state of the data" do not distinguish the claims from the prior art.

Appellant's arguments on page 15 of the remarks that "none of the situations taught or suggested by Arimilli describes a second node receives data from a first node having data in a modified state and assigning a shared state to the data at the second node, consistent" are clearly erroneous.

Examiner respectfully disagrees. As shown above, Arimilli clearly discloses "master 26 of a first agent, for example, processor complex 10a (i.e., first node), issues

modifying transaction 150a on system bus 12 that targets a cache line that is indicated as shared in the cache directory 22 of at least one agent 10 (i.e., second node); a result of a store request by the processor 16 in the second agent 10 that targets a cache line marked as shared in the second agent's cache processor 22; par. [0029]."

In support of such contention, appellant assumes in *Arimilli* "a target cache line is either: (a) already in the shared state (thus the data is not provided by another node) or (b) received from another agent 10".

However, it is clear that such contention cannot be supported for *Arimilli* makes it abundantly clear "master 26 of a first agent, for example, processor complex 10a (i.e., first node), issues modifying transaction 150a on system bus 12 that targets a cache line that is indicated as shared in the cache directory 22 of at least one agent 10 (i.e., second node); a result of a store request by the processor 16 in the second agent 10 that targets a cache line marked as shared in the second agent's cache processor 22; par. [0029];" where it is readily apparent that a node issues a transaction targeting a cache line and a result of a store request by a processor that targets the cache line marked as shared in the second node.

Additionally, in response to applicant's arguments against the references individually, one cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986).

Appellant's arguments on pages 15-16 of the remarks that none of the cited sections of Chen, nor Chen more generally, teach or suggest that "a first node is operative to provide the data to the second node and transition the associated state of the data at the first node from the modified state to an owner state in response to a non-migratory source broadcast request provided by the second node" recited in claim 1 are clearly erroneous.

First of all, in response to applicant's arguments against the references individually, one cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986).

Furthermore, Examiner respectfully disagrees with such contention. As described on pages 5, 8, of applicants' disclosure "a protocol employs a dirty (D-state) and an owner-shared state (O-state) to enable read migration without updating memory; the inclusion of the D-state and O-state in the cache coherency protocol allows for migration of data, without write-back, when requested from a cache containing the data in the M-state". Further, page 8, paragraph [0037] to page 9, paragraph [0039] provide instances where "a first node in response to a non-migratory source broadcast request provides data to a second node and transition an associated state of the data at a first node from modified state to an owner state without updating memory", namely:

[0037] Upon receiving the data from the target processor 14, the source processor 12 places the data in the appropriate cache line in its cache 22. The source processor 12 transitions the cache line of the cache 22 from the I-state to the D-state. In the D-state, the source processor 12 has the most up-to-date version of the data stored in its cache 22 and has not modified the data (e.g., by writing to the cache line in the cache 22). The data has thus migrated from the target processor 14 to the source processor 12 without write-back to (updating) the memory 16. Additionally, the D-state defines the processor 12 as a new cache ordering point for the data.

[0038] Continuing with the above example, assume that a processor (e.g., processor 14) broadcasts a migratory read request (XREADM) for the data, which is stored in the D-state in cache 22 of processor 12, as may occur after the initial read migration described above. The source processor 14 broadcasts an XREADM request to all other processors in the system 10, including the target processor 12 and those located in the other nodes 20, as well as to the memory 16. If the target processor 12 has not modified (e.g., written) the data, the target processor responds by providing a shared data response to the source processor 14 and transitioning the target processor cache line associated with the data from the D-state to the O-state. The requested data is received at the source processor 14 and placed in the appropriate cache line in the source

processor cache 24. The source processor 14 cache line transitions from the I-state to the S-state because, at this point, an up-to-date copy of the data is shared by both the source processor 14 (S-state) and the target processor 12 (O-state). Migration from the target processor 12 to the source processor 14 has not occurred because the target had not modified the data at the time of the XREADM request.

[0039] In the S-state, the source processor 14 has a valid and unmodified copy of the data. Since other processors may have valid copies of the data, the source processor 14 (being in the S-state) cannot respond to snoops by returning data and cannot write-back the data to memory 16. In the O-state, the target processor 12 has the most up-to-date version of the data stored in its cache 22. The target processor 12 cannot modify the data, and must write-back the data to memory 16 upon displacement (e.g., upon a write request or invalidate issued by another node). The target processor 12, being in the O-state, can respond to read requests by returning shared data, and thus may respond to subsequent read requests from other processors.

Similarly, Chen unequivocally discloses "the data in the cache has been modified and exclusively owned by a certain processor of the local node, and thus has become different from that existing in the home node; a command is issued to read an exclusive copy of the specific data which is stored in the local node and has not been modified, or read an exclusive

copy of the specific data which is stored in the local node and has been modified, or modify a shared copy of the specific data and exclusively own the specific data; changing the state of the local memory line of the remote node where the specific data is stored, from GONE into HOME, or changing the state of the local memory line of the remote node, from SHARED to HOME; col. 5, ll 5-10, 36-42; col. 5, ll 60 to col. 6, ll 6." Chen further discloses "discriminating whether the data has been modified by a first certain one of the at least two processors on the condition that a cache line of the external cache, where the data is stored, is in a CLEAN state; changing the state of the cache line from the CLEAN state to either of DIRTY-SHARED and DIRTY-ONLY state if the data has been modified into a modified data and allowing a second certain one of the at least two processors to directly request the modified data via a bus inside the first node when the cache line is in the DIRTY SHARED state; col. 3, ll 55-65; In CLEAN state, when the processor 112 is going to read and modify the exclusive copy of the data in the cache line, if the data is in the cache 1141 but has been modified, the modified data in the cache line, which becomes different from the data stored in the local memory 1231 of the home node 12, will be directly read off the bus 115 and exclusively owned by the processor 112, and thus the state of the cache line changes from CLEAN to DIRTY-ONLY; in CLEAN state, when the processor 112 is going to get a shared copy of the data in the cache line, and permits

another processor to share the data...if the data is in the cache 1141 but has been modified...the modified data in the cache line, which becomes different from the data stored in the local memory 1231 of the home node 12, will be directly read off bus 115 by the processor 112 and shared with other processors, and thus the state of the cache line changes from CLEAN to DIRY-SHARED; col. 6, ll 48 to col. 7, ll 2."

Thus, Chen clearly discloses, in the manner described in applicants' disclosure, "a protocol employs a dirty (D-state) and an owner-shared state (O-state) to enable read migration without updating memory; the inclusion of the D-state and O-state in the cache coherency protocol allows for migration of data, without write-back, when requested from a cache containing the data in the M-state, and as such discloses "a first node in response to a non-migratory source broadcast request provides data to a second node and transition an associated state of the data at a first node from modified state to an owner state without updating memory; col. 3, ll 55-65; col. 5, ll 5-10, 36-42; col. 5, ll 60 to col. 6, ll 6; col. 6, ll 48 to col. 7, ll 2".

Appellant's arguments on page 16 that "since no non-migratory source broadcast request exists in Chen, there is no basis to conclude that the teachings of Chen would be operative to provide data to the second node in response to such non-migratory

broadcast request. Therefore, Arimilli taken in view of Chen fails to teach or suggest that a first node is operative, in response to a non-migratory source broadcast request, to provide data to a second node" are clearly erroneous.

As shown above, Chen clearly discloses, in the manner described in applicants' disclosure, "a protocol employs a dirty (D-state) and an owner-shared state (O-state) to enable read migration without updating memory; the inclusion of the D-state and O-state in the cache coherency protocol allows for migration of data, without write-back, when requested from a cache containing the data in the M-state, and as such discloses "a first node in response to a non-migratory source broadcast request provides data to a second node and transition an associated state of the data at a first node from modified state to an owner state without updating memory; col. 3, ll 55-65; col. 5, ll 5-10, 36-42; col. 5, ll 60 to col. 6, ll 6; col. 6, ll 48 to col. 7, ll 2".

Additionally, Examiner would like to point out that while these limitations require "a second node" and "a first node", the functional recitations: "a node operative to [or capable of] receiving data" and "assigning a shared state to an associated state of the data" do not distinguish the claims from the prior art. While features of a system may be recited either structurally or functionally, claims directed to a system must be distinguished from the prior art in terms of structure rather than function. *In re Schreiber*, 128 F.3d 1473, 1477-78, 44 USPQ2d 1429, 1431-32 (Fed. Cir. 1997) (The absence of a disclosure in a prior art reference relating to function did not defeat the Board's finding of anticipation

of claimed apparatus because the limitations at issue were found to be inherent in the prior art reference); see also *In re Swinehart*, 439 F.2d 210, 212-13, 169 USPQ 226, 228-29 (CCPA 1971); < *In re Danly*, 263 F.2d 844, 847, 120 USPQ 528, 531 (CCPA 1959). "System claims cover what a device *is*, not what a device *does*." *Hewlett-Packard Co. v. Bausch & Lomb Inc.*, 909 F.2d 1464, 1469, 15 USPQ2d 1525, 1528 (Fed. Cir. 1990) (emphasis in original).

Appellant's arguments on pages 16-17 that Chen fails to teach or suggest "in response to the non-migratory source broadcast request, a first node transitions the state of the data at such node from a modified state to an owner state without updating system memory" recited in claim 1 are clearly erroneous.

Examiner strongly disagrees. Chen discloses, at least in Figs 2-5 how cache states are transitioned where it is readily apparent that such transitions result from data requests and responses of another node; "data in cache line is modified by certain processor; other processor may share modified data; another processor directly requests modified data; modified data is exclusively owned by certain processor; data in a cache line is modified, and the cache line changes from 1st to 2nd status; local memory line is in transition status; cache line recovers from 2nd to 1st status; Fig.5."

Additionally, in response to applicant's arguments against the references individually, one cannot show nonobviousness by attacking references

individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986).

Appellant's arguments on pages 17-18 that there is no motivation to combine Arimilli and Chen and that Chen teaches away from Arimilli are clearly erroneous.

In response to applicant's argument that there is no teaching, suggestion, or motivation to combine the references, the examiner recognizes that obviousness may be established by combining or modifying the teachings of the prior art to produce the claimed invention where there is some teaching, suggestion, or motivation to do so found either in the references themselves or in the knowledge generally available to one of ordinary skill in the art. See *In re Fine*, 837 F.2d 1071, 5 USPQ2d 1596 (Fed. Cir. 1988), *In re Jones*, 958 F.2d 347, 21 USPQ2d 1941 (Fed. Cir. 1992), and *KSR International Co. v. Teleflex, Inc.*, 550 U.S. 398, 82 USPQ2d 1385 (2007).

Appellant asserts on page 18 of the remarks that one of ordinary skill in the art would not combine and modify the teachings of Chen comprising a DSM system with the system of Arimilli comprising a central memory system because the person of

ordinary skill would recognize that the transactions required to maintain coherency in a DSM system would adversely affect the principal operation of a central memory system.

However, such assumptions are just mere allegations for there is nothing either in *Arimilli* or in *Chen* supporting appellant's position and combining *Arimilli* and *Chen* would provide no adverse result. Finally, , in determining obviousness under 35 U.S.C 103 in view of the Supreme Court decision in *KSR International Co. v. Teleflex Inc.*, the Supreme Court stated for purposes of 35 U.S.C 103, prior art can be either in the field of applicant's endeavor or be reasonably pertinent to the particular problem with which the applicant was concerned. Furthermore, **prior art that is in a field of endeavor other than that of the applicant, or solves a problem which is different from that which the applicant was trying to solve, may also be considered for the purposes of 35 U.S.C 103.** See, e.g., *In re KSR International Co. v. Teleflex Inc.*, 550 U.S. at __, 82 USPQ2d at 1396 (2007).

In response to applicant's argument that "the transactions required to maintain coherency in a DSM system would adversely affect the principal operation of a central memory system", the test for obviousness is not whether the features of a secondary reference may be bodily incorporated into the structure of the primary reference; nor is it that the claimed invention must be expressly suggested in any one or all of the references. Rather, the test is what the combined teachings of the references would have suggested to those of

ordinary skill in the art. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981).

Appellant's arguments on page 19, paragraph 3 that Chen fails to teach or suggest that "a first node is operative, in response to a non-migratory source broadcast request, to provide a shared data response to a second node" recited in claim 8 are clearly erroneous.

As shown above, Chen clearly discloses, in the manner described in applicants' disclosure, "a protocol employing a dirty (D-state) and an owner-shared state (O-state) to enable read migration without updating memory; the inclusion of the D-state and O-state in the cache coherency protocol allows for migration of data, without write-back (update), when requested from a cache containing the data in the M-state", and as such discloses "a first node being operative to respond to a non-migratory source broadcast; col. 3, ll 55-65; col. 5, ll 5-10, 36-42; col. 5, ll 60 to col. 6, ll 6; col. 6, ll 48 to col. 7, ll 2".

Additionally, Examiner would like to point out that while these limitations require "a first node", the functional recitation: "a node operative to [or capable of] responding to a non-migratory source broadcast " do not distinguish the claims from the prior art. While features of a system may be recited either structurally or functionally, claims directed to a system must be distinguished from the prior art in terms of structure rather than function. *In re Schreiber*, 128

F.3d 1473, 1477-78, 44 USPQ2d 1429, 1431-32 (Fed. Cir. 1997) (The absence of a disclosure in a prior art reference relating to function did not defeat the Board's finding of anticipation of claimed apparatus because the limitations at issue were found to be inherent in the prior art reference); see also *In re Swinehart*, 439 F.2d 210, 212-13, 169 USPQ 226, 228-29 (CCPA 1971); < *In re Danly*, 263 F.2d 844, 847, 120 USPQ 528, 531 (CCPA 1959). "System claims cover what a device *is*, not what a device *does*." *Hewlett-Packard Co. v. Bausch & Lomb Inc.*, 909 F.2d 1464, 1469, 15 USPQ2d 1525, 1528 (Fed. Cir. 1990) (emphasis in original).

Additionally, in response to applicant's arguments against the references individually, one cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986).

Appellant's argument on page 20, paragraph 3 that Arimilli in view of Chen fails to teach or suggest that "data (requested by a non-migratory source broadcast read request) is stored in a second processor node cache line, wherein an associated state of the second processor node cache line is assigned a shared state" recited in claim 12 is clearly erroneous.

Examiner strongly disagrees. Chen discloses, at least in Figs 2-5 how cache states are transitioned where it is readily apparent that such transitions result from data requests and responses of another node; "data in cache line is modified by certain processor (e.g., second node); other processor (e.g., second node) may share modified data; another processor directly requests modified data; modified data is exclusively owned by certain processor; data in a cache line is modified, and the cache line changes from 1st to 2nd status; Fig.5."

Appellant's argument on page 21 that Arimilli in view of Chen fails to teach or suggest "a first processor node that is programmed to respond to a non-migratory source broadcast read request of the second processor node by providing a shared data response to the second processor node" recited in claim 12 is clearly erroneous.

As shown above, with respect to claim 1, Chen clearly discloses, in the manner described in applicants' disclosure, "a protocol employs a dirty (D-state) and an owner-shared state (O-state) to enable read migration without updating memory; the inclusion of the D-state and O-state in the cache coherency protocol allows for migration of data, without write-back, when requested from a cache containing the data in the M-state", and as such discloses "a first processor node that is programmed to respond to a non-migratory source broadcast read request of the second processor node by providing a shared data response to the second

processor node; col. 3, ll 55-65; col. 5, ll 5-10, 36-42; col. 5, ll 60 to col. 6, ll 6;
col. 6, ll 48 to col. 7, ll 2".

Appellant's argument on page 21, paragraph 1 that Chen fails to teach that any response (from a first node) to a non-migratory source broadcast read request can be provided to a second node" allegedly recited in claim 12 is clearly erroneous.

As shown above, Chen clearly discloses, in the manner described in applicants' disclosure, "a protocol employing a dirty (D-state) and an owner-shared state (O-state) to enable read migration without updating memory; the inclusion of the D-state and O-state in the cache coherency protocol allows for migration of data, without write-back (Update), when requested from a cache containing the data in the M-state", and as such discloses "non-migratory source broadcast request issued by a source processor node; col. 3, ll 55-65; col. 5, ll 5-10, 36-42; col. 5, ll 60 to col. 6, ll 6; col. 6, ll 48 to col. 7, ll 2".

Additionally, in response to applicant's arguments against the references individually, one cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986).

Appellant's argument on page 21 that Arimilli in view of Chen fails to teach or suggest "transitioning an associated state of the first processor from a modified state to an owner state without updating the memory" recited in claim 12 is clearly erroneous.

As shown above, Chen clearly discloses, in the manner described in applicants' disclosure, "a protocol employing a dirty (D-state) and an owner-shared state (O-state) to enable read migration without updating memory; the inclusion of the D-state and O-state in the cache coherency protocol allows for migration of data, without write-back (Update), when requested from a cache containing the data in the M-state", and as such discloses "transitioning an associated state of the first processor from a modified state to an owner state without updating the memory; col. 3, ll 55-65; col. 5, ll 5-10, 36-42; col. 5, ll 60 to col. 6, ll 6; col. 6, ll 48 to col. 7, ll 2".

At least in Figs 2-5, Chen shows how cache states are transitioned where it is readily apparent that such transitions result from data requests and responses of another node; "data in cache line is modified by certain processor; other processor may share modified data; another processor directly requests modified data; modified data is exclusively owned by certain processor; data in a cache line is modified, and the cache line changes from 1st to 2nd status; local memory line is in transition status; cache line recovers from 2nd to 1st status; Fig.5."

Appellant's argument on page 21, paragraph 4 to page 22 that Chen fails to teach or suggest "any command that enables a second processor node to receive data from a processor node with the data in either of the DIRTY-ONLY and DIRTY-SHARED states indicate that data in an L3 cache has been modified" is clearly erroneous.

In response to applicant's argument that the references fail to show certain features of applicant's invention, it is noted that the features upon which applicant relies (i.e., any command that enables a second processor node to receive data from a processor node with the data in either of the DIRTY-ONLY and DIRTY-SHARED states indicate that data in an L3 cache has been modified) are not recited in the rejected claim(s). Although the claims are interpreted in light of the specification, limitations from the specification are not read into the claims. See *In re Van Geuns*, 988 F.2d 1181, 26 USPQ2d 1057 (Fed. Cir. 1993).

Additionally, in response to applicant's arguments against the references individually, one cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986).

Appellant's argument on page 2 paragraph 2 to page 23 that "there is not proper motivation to combine Arimilli and Chen since Chen relates to management of an L3

cache by DSM controller whereas claim 12 recites first and second processor nodes having cache lines and associated cache states" is in error.

However, there is nothing either in *Arimilli* or in *Chen* supporting appellant's position. Additionally, in determining obviousness under 35 U.S.C 103 in view of the Supreme Court decision in *KSR International Co. v. Teleflex Inc.*, the Supreme Court stated for purposes of 35 U.S.C 103, prior art can be either in the field of applicant's endeavor or be reasonably pertinent to the particular problem with which the applicant was concerned. Furthermore, prior art that is in a field of endeavor other than that of the applicant, or solves a problem which is different from that which the applicant was trying to solve, may also be considered for the purposes of 35 U.S.C 103. See, e.g., *In re KSR International Co. v. Teleflex Inc.*, 550 U.S. at ___, 82 USPQ2d at 1396 (2007).

Appellant's argument on page 24 that *Chen* fails to teach or suggest a relationship between three nodes (or processor nodes, as recited in claim 19): first second and at least one other node - and also fails to teach or suggest a system with source broadcast requests" is clearly in error.

Examiner would like to point out that Applicant is reading limitations of the specification into the claims to thereby narrow the scope of the claims by implicitly adding disclosed limitations which have no express basis in the claims. This is impermissible importation of subject matter from the specification into the claim and such is not in accordance with USPTO rules and procedures. See

MPEP 2111. See also *In re Morris*, 127 F.3d 1048, 1054-55, 44 USPQ2d 1023, 1027-28 (Fed. Cir. 1997).

Additionally, Chen clearly discloses “a system including at least (i.e., 1-n processor nodes) a first and a second nodes including cache for storing data and at least two processors wherein the data can be modified by at least one of the at least two processors and a node may directly request data from another node (i.e., broadcast request); Abstract”.

Furthermore, in response to applicant’s arguments against the references individually, one cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986).

Applicant’s argument on page 25 that Arimilli taken in view of Chen fails to teach or suggest “data (requested by a non-migratory source broadcast read request (XREADN)) is provided to a first node, wherein a shared state is associated with the data at the first node in response to the first node receiving the data from a second node” recited in claim 24 is clearly erroneous.

Examiner strongly disagrees with such contention. As described on pages 5, 8, of applicants’ disclosure “a protocol employs a dirty (D-state) and an owner-shared state (O-state) to enable read migration without updating memory; the

inclusion of the D-state and O-state in the cache coherency protocol allows for migration of data, without write-back, when requested from a cache containing the data in the M-state". Further, page 8, paragraph [0037] to page 9, paragraph [0039] provide instances where "a first node in response to a non-migratory source broadcast request provides data to a second node and transition an associated state of the data at a first node from modified state to an owner state without updating memory", namely:

[0037] Upon receiving the data from the target processor 14, the source processor 12 places the data in the appropriate cache line in its cache 22. The source processor 12 transitions the cache line of the cache 22 from the I-state to the D-state. In the D-state, the source processor 12 has the most up-to-date version of the data stored in its cache 22 and has not modified the data (e.g., by writing to the cache line in the cache 22). The data has thus migrated from the target processor 14 to the source processor 12 without write-back to (updating) the memory 16. Additionally, the D-state defines the processor 12 as a new cache ordering point for the data.

[0038] Continuing with the above example, assume that a processor (e.g., processor 14) broadcasts a migratory read request (XREADM) for the data, which is stored in the D-state in cache 22 of processor 12, as may occur after the initial read migration described above. The source processor 14 broadcasts an XREADM request to all other processors in

the system 10, including the target processor 12 and those located in the other nodes 20, as well as to the memory 16. If the target processor 12 has not modified (e.g., written) the data, the target processor responds by providing a shared data response to the source processor 14 and transitioning the target processor cache line associated with the data from the D-state to the O-state. The requested data is received at the source processor 14 and placed in the appropriate cache line in the source processor cache 24. The source processor 14 cache line transitions from the I-state to the S-state because, at this point, an up-to-date copy of the data is shared by both the source processor 14 (S-state) and the target processor 12 (O-state). Migration from the target processor 12 to the source processor 14 has not occurred because the target had not modified the data at the time of the XREADM request.

[0039] In the S-state, the source processor 14 has a valid and unmodified copy of the data. Since other processors may have valid copies of the data, the source processor 14 (being in the S-state) cannot respond to snoops by returning data and cannot write-back the data to memory 16. In the O-state, the target processor 12 has the most up-to-date version of the data stored in its cache 22. The target processor 12 cannot modify the data, and must write-back the data to memory 16 upon displacement (e.g., upon a write request or invalidate issued by another node). The target processor 12, being in the O-state, can respond to read requests by

returning shared data, and thus may respond to subsequent read requests from other processors.

Similarly, Chen unequivocally discloses "the data in the cache has been modified and exclusively owned by a certain processor of the local node, and thus has become different from that existing in the home node; a command is issued to read an exclusive copy of the specific data which is stored in the local node and has not been modified, or read an exclusive copy of the specific data which is stored in the local node and has been modified, or modify a shared copy of the specific data and exclusively own the specific data; changing the state of the local memory line of the remote node where the specific data is stored, from GONE into HOME, or changing the state of the local memory line of the remote node, from SHARED to HOME; col. 5, ll 5-10, 36-42; col. 5, ll 60 to col. 6, ll 6." Chen further discloses "discriminating whether the data has been modified by a first certain one of the at least two processors on the condition that a cache line of the external cache, where the data is stored, is in a CLEAN state; changing the state of the cache line from the CLEAN state to either of DIRTY-SHARED and DIRTY-ONLY state if the data has been modified into a modified data and allowing a second certain one of the at least two processors to directly request the modified data via a bus inside the first node when the cache line is in the DIRTY_SHARED state; col. 3, ll 55-65; In CLEAN state, when the processor 112 is going to read and modify the exclusive copy of the data in the cache line, if the data is in the cache 1141 but has been modified, the modified data in the cache

line, which becomes different from the data stored in the local memory 1231 of the home node 12, will be directly read off the bus 115 and exclusively owned by the processor 112, and thus the state of the cache line changes from CLEAN to DIRTY-ONLY; in CLEAN state, when the processor 112 is going to get a shared copy of the data in the cache line, and permits another processor to share the data...if the data is in the cache 1141 but has been modified...the modified data in the cache line, which becomes different from the data stored in the local memory 1231 of the home node 12, will be directly read off bus 115 by the processor 112 and shared with other processors, and thus the state of the cache line changes from CLEAN to DIRTY-SHARED; col. 6, ll 48 to col. 7, ll 2."

Thus, Chen clearly discloses, in the manner described in applicants' disclosure, "a protocol employs a dirty (D-state) and an owner-shared state (O-state) to enable read migration without updating memory; the inclusion of the D-state and O-state in the cache coherency protocol allows for migration of data, without write-back, when requested from a cache containing the data in the M-state, and as such discloses "a first node in response to a non-migratory source broadcast request provides data to a second node and transition an associated state of the data at a first node from modified state to an owner state without updating memory; col. 3, ll 55-65; col. 5, ll 5-10, 36-42; col. 5, ll 60 to col. 6, ll 6; col. 6, ll 48 to col. 7, ll 2".

Appellant's further argue on page 26, paragraph 2 that Arimilli in view of Chen fails to teach or suggest "means for providing data from a second node in response to a XREADN request" recited in claim 24.

As shown above, Chen discloses "providing data from a second node in response to a XREADN request". It is worth mentioning that the "means for providing" is a processor as shown on page 12, line 8 of the Brief files November 30, 2007, and the combination of Arimilli and Chen clearly discloses such a processor.

Appellant's argument on page 26, paragraph 3 that Arimilli in view of Chen fails to teach or suggest "means for transitioning a modified state associated with data at a second node to an owner state without updating the memory" recited in claim 24, is clearly erroneous.

As shown above with respect to claim 1, the combination of Arimilli and Chen discloses "transitioning a modified state associated with data at a second node to an owner state without updating the memory". It is worth mentioning that the "means for transitioning a modified state" is a processor as identified on page 12 of the Brief filed on November 30, 2007, and the combination of Arimilli and Chen discloses such a processor.

Appellant's argument on page 28 with respect to claim 25 has been addressed in the rejection of the claims and the paragraphs above and the same reasoning applies to limitations similar to those of claim 1 as shown above.

Appellant's argument on page 29 with respect to claim 28 has been addressed in the rejection of the claims and the paragraphs above and the same reasoning applies to limitations similar to those of claim 1 as shown above.

Appellant's arguments on page 31, paragraph 2 regarding claims 30 and 31 that Chen does not teach or suggest anywhere "the selection between a non-migratory request and a migratory request to rebroadcast from the first node."

In response to applicant's argument that the references fail to show certain features of applicant's invention, it is noted that the features upon which applicant relies (i.e., the selection between a non-migratory request and a migratory request to rebroadcast from the first node) are not recited in the rejected claim(s). Although the claims are interpreted in light of the specification, limitations from the specification are not read into the claims. See *In re Van Geuns*, 988 F.2d 1181, 26 USPQ2d 1057 (Fed. Cir. 1993).

Appellant's argument on page 32 with respect to claim 32 has been addressed in the rejection of the claims above and in the previous paragraphs

where same reasoning applies to limitations similar to those of claim 1 as shown above.

Appellant's argument on page 34 with respect to claim 21 has been addressed in the rejection of the claims above and in the paragraph supra where the same reasoning applies to limitations similar to those of claim 1 as shown above.

In view of the foregoing, it is manifest that the claimed invention is not patentably distinct over the combinations of Arimilli, Chen, and Cypher when interpreted in light of the specification and taken into account legal standards for determining obviousness under 35 U.S.C. 103 in view of the Supreme Court decision *In re KSR International Co. v. Teleflex Inc.*, 550 U.S. at __, 82 USPQ2d at 1396 (2007).

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,

April 23, 2010

/Mardochee Chery/

Primary Examiner, Art Unit 2188

Conferees:

/Matt Kim/

Supervisory Patent Examiner, Art Unit 2186

/Kevin L Ellis/

Supervisory Patent Examiner, Art Unit 2117